

HOMOTOPY LIMITS IN COQ

JEREMY AVIGAD, KRZYSZTOF KAPULKIN, AND PETER LEFANU LUMSDAINE

ABSTRACT. Working in Homotopy Type Theory, we provide a systematic study of basic homotopy limits and related constructions. The entire development has been formally verified in the Coq interactive proof assistant.

CONTENTS

Introduction	1
1. The logical setting	3
1.1. The theory under consideration	3
1.2. Semantics	4
1.3. Notation and terminology	5
2. Fundamental constructions	7
2.1. Background from homotopy type theory	7
2.2. Fibration categories from type theory	9
3. Limits and applications	14
3.1. Pullbacks and equalizers	14
3.2. Limits	19
3.3. Pointed types and fiber sequences	22
4. Reflections on the formal verification	23
4.1. Limitations	23
4.2. Proof-relevance	23
4.3. Constructing paths	24
4.4. General strategies	26
4.5. A case study: the two pullbacks lemma	27
References	29

INTRODUCTION

Homotopy type theory is based on the discovery that formal dependent type theory has a natural homotopy-theoretic interpretation ([Voe06], [AW09]). Since a number of interactive proof assistants implement versions of dependent type theory, these observations open the possibility of developing parts of homotopy theory formally with the help of such assistants; see [PW12] for a helpful overview.

Date: March 1, 2013.

In this spirit, we carry out a number of homotopy-theoretic constructions in a core system of homotopy type theory. In particular, we define and investigate (homotopy) pullbacks, equalizers, limits over graphs, pointed spaces, and fiber sequences. The entire development is formalized with the Coq interactive proof assistant.

Besides the formalization itself, we also compare the semantics of type theory with fibration categories, a standard homotopy-theoretic setting for the construction of homotopy limits.

We should mention that many of the facts we present below are already known in folklore; in any case, none of them will be unexpected to researchers in the field. Egbert Rijke and Bas Spitters [RS13] have also independently investigated limits and colimits over graphs within a similar type theory. We hope it will prove useful, however, to have a systematic treatment of these basic results, fully formalized in Coq and available as a library for future use. We also hope that the practical lessons we learned during the formalization process may be useful to others.

Our Coq development builds on a library for homotopy type theory developed jointly by various people, under the leadership of Andrej Bauer, Lumsdaine, and Michael Shulman [HoT]. Another extensive library has been developed by Vladimir Voevodsky [Voe], and some of our verified results overlap his.

Outline. In Section 1, we set out the formal framework of our work: the type theory under consideration, and its intended interpretation. In Section 2, we review some background on the type-theoretic development of homotopy theory, and use it to show that every categorical model of the theory carries the structure of a fibration category. Section 3 presents the main body of our formalization: a concise treatment of the content, in traditional mathematical prose. Finally, in Section 4 we share some reflections on practical aspects of the formalization process.

Our formal development in Coq can be found online at

<https://github.com/peterlefanulumsdaine/hott-limits/>.

Acknowledgements. Most of this work was carried out during the Special Year on Univalent Foundations at the Institute for Advanced Study, and all three authors are grateful to the Institute for its hospitality and support.

Avigad’s work has been partially supported by NSF grant DMS-1068829 and AFOSR grant FA9550-12-1-0370. Kapulkin was supported by NSF Grant DMS-1001191 (P.I. Steve Awodey) and a grant from the Benter Foundation (P.I. Thomas Hales). Lumsdaine was supported by NSF grant DMS-1128155. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Kapulkin dedicates this work to his mother.

1. THE LOGICAL SETTING

In this section, we first lay out the specific logical system in which we will work; we then review its intended semantics, insofar as they are relevant to working within the theory; finally, we fix the basic notation and terminology we will use, based on the intended semantics.

1.1. The theory under consideration. We assume familiarity with predicative dependent type theory, often known as Martin-Löf type theory [ML84]. The following types, and associated rules, form a minimal core to that system:

- (1) dependent products $\Pi_{x:A}B$, and the associated introduction, elimination, and computation rules
- (2) dependent sums $\Sigma_{x:A}B$, and the associated introduction, elimination, and computation rules
- (3) identity types Id_A , and the associated introduction, elimination, and computation rules

Most developments in homotopy type theory add at least the following rule:

- (4) function extensionality: for any type A , any type B depending on $x : A$, and functions $f, g : \Pi_{x:A}B$, if $fx = gx$ for every $x : A$, then $f = g$.

The system based on these rules is used in [AGS12], where it is denoted by \mathcal{H} ; it also forms a sufficient basis for much of the present formalization. Some constructions, in addition, depend on:

- (5) the type Nat of natural numbers, with the usual introduction, elimination, and computation rules,

from which the empty type, unit type, and other finite types can be defined; and finally, some of our definitions presuppose the existence of:

- (6) a universe U of types, containing Nat , and closed under the formation of dependent products, sums, and identity types.

We use quantification over the universe to define the universal properties of pullbacks and limits, but also give equivalent formulations that do not make use of such a universe.

Besides these, the version of Coq we used implements η -conversion for functions, $\lambda x. fx = f$, as a built-in conversion rule. As a propositional equality, it is derivable from function extensionality, so we do not believe its use is essential; however, since it is unavoidably present in the proof assistant, we include it in our formal theory.

In sum, if we take axioms (1)–(3) to represent the core of Martin-Löf type theory, ML , it is then reasonable to denote our overall framework as

$$\text{ML} + (\text{funext}) + (\eta) + (\text{Nat}) + (\text{U}).$$

For brevity, we refer to this in the present paper as \mathcal{H}' ; thus the mathematical content of our work is that the constructions and assertions of Sections 2

and **3** are consequences of that formal theory. As noted above, however, most of our results do not require **Nat**, and many do not require **U**.

We do not consider in the present work extra axioms such as Univalence, resizing, or higher inductive types.

One final comment about the formal verification: rather than providing **Id**, **Nat**, **Bool**, and so on individually, Coq provides a general mechanism for defining inductive types, which these are then defined as instances of. However, the resulting eliminators for these types correspond precisely to the rules for them described above. Coq also provides (dependent) record types, as syntactic sugar for certain inductive types; in some cases, using record types made type checking more efficient, and brought notational benefits as well. As these may be routinely translated into (iterated) Σ -types, their use has no bearing on the question of derivability in \mathcal{H}' .

1.2. Semantics.

1.2.1. *General algebraic semantics.* The fully general semantics of dependent type theories are, from a purely algebraic point of view, well-understood. Essentially, a model of a dependent type theory \mathbf{T} with the same basic judgements and structural rules as \mathcal{H}' may be defined as a *contextual category*—that is, a category equipped with structure sufficient to model the structural rules of the theory—along with further algebraic structure corresponding to the particular logical constructors under consideration. For the details of this definition, see [Str91]; for brevity, we will refer to such a structure as a *categorical model of \mathbf{T}* .

The justification for calling such structures models comes from the fact that the syntax of the theory forms an initial such structure:

Definition 1.2.1. Given any dependent type theory \mathbf{T} , the *syntactic category* $\mathcal{C}(\mathbf{T})$ is given as follows:

- objects of $\mathcal{C}(\mathbf{T})$ are contexts $[x_1:A_1, \dots, x_n:A_n]$ of \mathbf{T} , up to definitional equality and renaming of free variables;
- maps of $\mathcal{C}(\mathbf{T})$ are *context morphisms* (a.k.a. *substitutions*), again up to definitional equality and renaming of free variables. That is, a map

$$f : [x_1:A_1, \dots, x_n:A_n] \longrightarrow [y_1:B_1, \dots, y_m:B_m(y_1, \dots, y_{m-1})]$$

is represented by a sequence of terms

$$\begin{aligned} x_1:A_1, \dots, x_n:A_n &\vdash f_1 : B_1 \\ &\vdots \\ x_1:A_1, \dots, x_n:A_n &\vdash f_m : B_m(f_1, \dots, f_{m-1}). \end{aligned}$$

Moreover, $\mathcal{C}(\mathbf{T})$ may naturally be given the structure of a contextual category; for each logical rule of \mathbf{T} , $\mathcal{C}(\mathbf{T})$ carries the corresponding algebraic structure.

Fact 1.2.2 ([Str91]¹). $\mathcal{C}(\mathbf{T})$ is initial among categorical models of \mathbf{T} .

Thus any other categorical model \mathbf{C} has a canonical structure-preserving functor from $\mathcal{C}(\mathbf{T})$ —that is, an interpretation function, interpreting the syntax of \mathbf{T} in \mathbf{C} .

1.2.2. Homotopical semantics. Homotopy type theory is based on the realization ([HS98], [AW09], [vdBG12], [Voe10]) that various homotopy-theoretic settings give natural examples of such categorical models. Very roughly, a type A denotes a space; a family $B(x)$ of types, indexed over A , denotes of fibration over A ; a term $t(x)$ of type B , with variable x of type A , denotes a continuous map from A to B , and so on.

The main motivating interpretation, for us, is the model in simplicial sets—one of the most well-studied models of spaces in homotopy theory. The full details of this interpretation are rather technical, so since we never require them, we omit them here; see [KLV12] for a complete presentation of the simplicial set model, and [Shu12] for more general related models. We sketch here just the main ingredients of the interpretation, insofar as they justify the intuition and terminology for working within the theory.

In this model, closed types (and, more generally, contexts) are interpreted as Kan complexes; dependent types, as Kan fibrations. Most type formers— Π -types, Σ -types, \mathbf{Nat} , etc.—are interpreted as in the more familiar topos logic: Π -types by the right adjoint to pullback, Σ -types by the left, \mathbf{Nat} by the natural numbers object, and so on.

The main novelty, however, is the interpretation of the identity type $\mathrm{Id}_A(x, y)$ with variables x and y from A . In set- and topos-theoretic models, one would interpret it as the diagonal map $A \longrightarrow A \times A$. However, in simplicial sets (and other homotopy-theoretic settings) this map is hardly ever a fibration. It can, however, be replaced by a fibration $P(A) \longrightarrow A \times A$, where $P(A)$ is the *path object* of A ; this is then used to interpret the identity type of A . Thinking of a simplicial set as a space, $P(A)$ represents the space of paths in A , with the fibration $P(A) \longrightarrow A \times A$ giving the indexing of paths over their endpoints. In particular contrast to the set-theoretic situation, for given $x, y : A$ the space $P(A)(x, y)$ of paths from x to y may be not merely a proposition, but a non-trivial space in its own right.

1.3. Notation and terminology. The homotopical interpretation has guided the various choices of notation and terminology we use in the presentation of our development. In particular, we use notation $p : (x \rightsquigarrow y)$ for the identity type to indicate that we would like to think of it as the type of paths from x to y . The Homotopy Type Theory library provides both the notations $\mathbf{x} = \mathbf{y}$ and $\mathbf{x} \rightsquigarrow \mathbf{y}$ for this type, and while working in Coq, we prefer to use the former rather than the latter. However, in the informal presentation below,

¹Unfortunately, to our knowledge, no general form of this result exists in the literature; it is shown for certain specific type theories in [Str91] and elsewhere, and its extension to other combinations of the standard rules (such as \mathcal{H}') is well-known in folklore.

we find it most natural to understand our constructions as constructions of paths, rather than equality proofs; and so we have settled on using the latter notation, and favor the word “path” over “equality.”

In other respects, however, we have found it more convenient to leave the homotopy-theoretic interpretation implicit. For example, the natural definitions of pullbacks, equalizers, and limits in type-theoretic notation turn out to characterize homotopy pullbacks, homotopy equalizers, and homotopy limits in the homotopy-theoretic interpretation. Having kept the notion of “path” prominent, sprinkling the word “homotopy” everywhere seemed to impose an unnecessary burden; thus, both in code and in prose we refer just to “pullbacks,” “equalizers,” and “limits.” (This is customary in higher category theory (see, e.g., [Lur09]), when one uses, for example, the word “limit” for an object that in strict terms is merely a homotopy limit.)

For the sake of readability, we will use standard mathematical terminology and notation in Sections 2 and 3, rather than attempting to adhere closely to the notation used in the Coq code. The following table lists some of the basic notions of our development, comparing the notations used in our presentation here with those used in the Coq formalization:

informal notion	mathematical notation	Coq notation
p is a path from x to y	$p : (x \rightsquigarrow y)$	<code>p : x = y</code>
the identity path at x	$\text{refl}(x)$	<code>idpath x</code>
the concatenation of p and q	$p \bullet q$	<code>p @ q</code>
the inverse of p	\bar{p}	<code>!p</code>
B is a fibration over A	$B \twoheadrightarrow A$	<code>B : A -> Type</code>
the total space of B over A	$\Sigma_{x:A} B(x)$	<code>{ x : A & B x }</code>
the dependent product of B over A	$\Pi_{x:A} B(x)$	<code>forall x : A, B x</code>
e is an equivalence from A to B	$e : A \simeq B$	<code>e : A <~> B</code>
the inverse of e	e^{-1}	<code>e~1</code>
a universe of small types	\mathcal{U}	<code>UU</code>
the natural numbers	\mathbb{N}	<code>nat</code>

As usual in homotopy type theory, we represent logic using *propositions-as-types*, with implication, conjunction, and universal and existential quantification interpreted in terms of function, product, Π -, and Σ -types respectively. Thus, for example, the functional extensionality axiom (Axiom 4 in Section 1.1 above), is formally a constant of type:

$$\text{funext} : \prod_{\substack{A:\text{Type} \\ B:A \rightarrow \text{Type}}} \prod_{f,g:\prod_{x:A} B(x)} \left(\prod_{x:A} (f x \rightsquigarrow g x) \right) \rightarrow (f \rightsquigarrow g).$$

Notice that Σ -types provide a useful way of “packaging” related pieces of data into a single type: to illustrate this, consider Definition 3.1.5 below. Formally, a cospan consists of types A , B , and C , and maps $f : A \rightarrow C$, $g : B \rightarrow C$. Given a type X , a cone over this cospan with vertex X consists

of maps $h : X \longrightarrow A$ and $k : B \longrightarrow C$, and a family of paths $(f(hx) \rightsquigarrow g(kx))$ for each x in X . In other words, such a cone is an element of the type

$$\sum_{\substack{h : X \rightarrow A \\ k : X \rightarrow B}} \prod_{x : X} (f(hx) \rightsquigarrow g(kx)).$$

Thus our formal definition in Coq reads as follows:

```
Definition cospan_cone {A B C : Type} (f : A -> C)
  (g : B -> C) (X : Type)
:= { h : (X -> A) & { k : (X -> B)
  & forall x, paths (f(h x)) (g(k x)) } }.
```

The curly braces around the arguments A , B , and C indicate that these are treated as *implicit arguments*. This means that the user may write just `cospan_cone f g X`, leaving the system to infer A , B , and C from the types of f and g . Sometimes one needs to turn this feature off, and specify such arguments; writing `@cospan_cone A B C f g X` tells Coq to expect all the arguments of `cospan_cone` to be given explicitly.

2. FUNDAMENTAL CONSTRUCTIONS

In this section and the next, we develop the basic theory of homotopy limits and related notions in \mathcal{H}' . We have already explained, in Section 1, how the basic ingredients are represented in the language of Coq, and complete details of the whole development can be found in the files comprising our formal verification. Especially in Section 3, therefore, we will generally only sketch most proofs, leaving out steps that are straightforward and routine (and even some that are not).

2.1. Background from homotopy type theory. Our formal work builds on the HoTT library [HoT] for homotopy theory developed by Bauer, Lumsdaine, Shulman, and others. We begin by summarizing some of the basic components of this library that are used throughout.

2.1.1. Operations on paths. Given any $x, y : X$, we write $p : (x \rightsquigarrow y)$ to denote the fact that p is a path from x to y . For every x , there is an “identity path” $\text{refl}(x) : (x \rightsquigarrow x)$. The central property characterizing the type of paths is its elimination principle, which says roughly that to construct an object of a type $C(x, y, p)$ depending on a path p from x to y , it suffices to construct an element of $C(x, x, \text{refl}(x))$, in which p has been “contracted” to an identity path.

Paths admit various operations familiar from homotopy theory and higher category theory. Any two paths $p : (x \rightsquigarrow y)$ and $q : (y \rightsquigarrow z)$ can be concatenated, yielding a path $p \bullet q : (x \rightsquigarrow z)$. Moreover, $\text{refl}(x) : (x \rightsquigarrow x)$ is a unit element for this operation, and every path admits an inverse $\bar{p} : (y \rightsquigarrow x)$. These operations satisfy the groupoid laws, but, as in homotopy theory, only up to a higher path. For example, we can find an inhabitant of the type

$(p \cdot \bar{p} \rightsquigarrow \text{refl}(x))$. In fact, every type, together with the tower of its paths, forms an ∞ -groupoid of some sort; (precise statements along these lines can be found in [vdBG11], [Lum09]).

Moreover, the maps between types respect the paths and the structure on them. That is: given any $p : (x \rightsquigarrow y)$ in X and $f : X \rightarrow Y$, we obtain a path $f[p] : (f(x) \rightsquigarrow f(y))$; and this is functorial, in the up-to-homotopy sense that there is, for example, an inhabitant of the type $(f[p \cdot q] \rightsquigarrow f[p] \cdot f[q])$.

2.1.2. Equivalences and h-levels. The notion of paths allows us to recover several familiar notions from algebraic topology.

We can, for example, say that a type X is *contractible* if there is some $x_0 : X$, and a function giving for each $x : X$ a path $(x \rightsquigarrow x_0)$. One can also construct the *homotopy fiber* of a map $f : X \rightarrow Y$ over an element $y : Y$ by:

$$\text{hfib}(f, y) := \sum_{x:X} (f(x) \rightsquigarrow y).$$

Given these we say that a map $f : X \rightarrow Y$ is an *equivalence* if for all $y : Y$ the homotopy fiber of f over y is contractible. The HoTT library provides many crucial results on equivalences. For example, a map is an equivalence if it has a two-sided inverse (up to homotopy), or alternatively two one-sided inverses. We will return to these points in more detail in Section 2.2.

Another notion that smoothly transfers from algebraic topology to HoTT is the notion of an n -type. Classically, an n -type is a space whose homotopy groups vanish above dimension n . In HoTT we define a similar hierarchy, although with the indexing shifted by 2. The reason for this shift is simply that, formally, natural numbers are easier to deal with than integers greater than or equal to -2 ; in particular, in an interactive theorem prover, one can make use of the standard inductively defined type of natural numbers.

Precisely, *h-levels* are defined by induction. A type X is of h-level 0 if it is contractible; and of h-level $n + 1$ if for all $x, y : X$, the type $(x \rightsquigarrow y)$ of paths from x to y is of h-level n . Thus in particular, types of h-level 1 may be considered as propositions, carrying no more information than the fact of being inhabited; and types of h-level 2 are (up-to-homotopy) discrete sets. We call such types *h-props* and *h-sets* respectively.

2.1.3. Functional extensionality. Given two types X and Y , the type $X \rightarrow Y$ of maps between them can be equipped with the notion of a path (or rather, a “homotopy”) in two different ways. First, for any $f, g : X \rightarrow Y$, one can form $(f \rightsquigarrow g)$, in the usual way. On the hand, one can also compare two functions pointwise, asking for an element of $\prod_{x:X} (f(x) \rightsquigarrow g(x))$; we call such a function h a *homotopy* from f to g , and write $h : f \Rightarrow g$.

Given any $p : (f \rightsquigarrow g)$, by the elimination principle for paths, we obtain an element of the type $f \sim g$. The functional extensionality axiom asserts that this assignment is an equivalence; that is, that given a pointwise homotopy between two maps, we can always find a path between them in the function

type inducing the original homotopy. More generally, functional extensionality asserts this equivalence between paths and homotopies in *dependent* function types $\prod_{x:A} B(x)$.

2.1.4. Dependent sums. The interaction between dependent sums and paths is crucial in our work. Let $B(x)$ be a type depending on $x : A$. It is easy to see that then a path $p : (a \rightsquigarrow a')$ in A induces an equivalence $p_! : B(a) \longrightarrow B(a')$, which we call *transport* between fibers. As everything before, this commutes appropriately with the operations on paths; for example, for any $p : (a \rightsquigarrow a')$ and $q : (a' \rightsquigarrow a'')$, and $b : B(a)$ we have $((p \cdot q)_! b \rightsquigarrow q_!(p_! b))$.

This also provides a means to construct paths between two elements of a Σ -type. Given a path $p : ((a, b) \rightsquigarrow (a', b'))$ in a $\sum_{x:A} B(x)$, we get a pair of paths: $p_1 : (a \rightsquigarrow a')$ and $p_2 : ((p_1)_! b \rightsquigarrow b')$; and conversely, given such a pair of paths, we can recover the original path p . This construction is ubiquitous in the formalization, since so many objects are defined using Σ -types; for more discussion, see Section 4.3 below.

2.2. Fibration categories from type theory. In this section, we show that any categorical model of \mathcal{H}' (including its syntactic category) satisfies the axioms of a fibration category, following the lines of results such as [GG08], [Lum11].

The result follows from a combination of internal reasoning—proving certain statements in the type theory—and external (meta-theoretic), showing how in models, the internal statements translate into the desired axioms. Since we will be switching back and forth frequently between these two different logical settings, we use *sans serif text* to distinguish the *internal reasoning* from the external. After this, we look at how some standard properties of fibration categories translate in terms of the type theory.

The internal parts of this section are formalized in the file `Fundamentals.v`.

We start by recalling the definition of a fibration category (for more on which, see [Bro73], [Bau89]):

Definition 2.2.1. A *fibration category* is a category \mathbf{C} together with two distinguished classes of maps, \mathcal{W} (*weak equivalences*) and \mathcal{F} (*fibrations*) satisfying the following conditions:

- (1) Weak equivalences satisfy 2-out-of-6 condition; i.e., given a composable triple of morphisms

$$W \xrightarrow{f} X \xrightarrow{g} Y \xrightarrow{h} Z,$$

if $g \cdot f$ and $h \cdot g$ are weak equivalences, then so are f , g , h , and $h \cdot g \cdot f$.

- (2) \mathcal{F} is closed under composition.
- (3) Calling a map that is both a weak equivalence and a fibration an *acyclic fibration*, all isomorphisms are acyclic fibrations.
- (4) \mathbf{C} has a terminal object $\mathbf{1}$.
- (5) Pullbacks along fibrations exist; fibrations and acyclic fibrations are stable under pullback.

- (6) For any object $X \in \mathbf{C}$, the diagonal morphism $\Delta: X \longrightarrow X \times X$ can be factored as a weak equivalence followed by a fibration.

$$X \longrightarrow PX \longrightarrow X \times X.$$

(Such a factorization, and by abuse of language also the object PX , is called a *path object* for X .)

- (7) Every object is *fibrant*; that is, the unique map $X \longrightarrow \mathbf{1}$ is a fibration, for any $X \in \mathbf{C}$.

Remark 2.2.2. This is slightly stronger than the original definition given by Brown, in that it requires the class \mathcal{W} to satisfy the 2-out-of-6 axiom rather than just the more familiar 2-out-of-3. However, once \mathbf{C} satisfies all the other axioms, the following conditions are equivalent (the result is due to Cisinski; see [RB06, Thm. 7.2.7]):

- (1) \mathcal{W} satisfies 2-out-of-6;
- (2) \mathcal{W} satisfies 2-out-of-3 and is *saturated*; that is, if a map w of \mathbf{C} becomes an isomorphism in $\text{Ho}(\mathbf{C})$, then $w \in \mathcal{W}$.

In this section we show that any categorical model of \mathcal{H}' (in the sense of Section 1.2) carries the structure of a fibration category; and so, in particular, the syntactic category $\mathcal{C}(\mathcal{H}')$ does. From here on, fix some categorical model \mathbf{C} of \mathcal{H}' .

For convenience of exposition, we also assume in this section strong η -rules for Σ -types, so every context is isomorphic to (a context consisting of just) a single iterated Σ -type: for instance,

$$[x:A, y:B(x)] \cong [p : \sum_{x:A} B(x)].$$

This allows us to work just with types, rather than with general contexts. However, nothing here depends on that assumption; one may simply replace types with contexts and Σ -types with context extensions, in particular in the definition of the fibrations:

Definition 2.2.3 (Gambino–Garner [GG08]). A map of \mathbf{C} is a *fibration* if it is isomorphic to some composite of first projections from Σ -types,

$$\sum_{x:A} B(x) \rightarrow A.$$

Denote the class of fibrations by \mathcal{F} .

(This is a slight simplification of Gambino and Garner’s original definition, which also closes \mathcal{F} under retracts.) Note that “isomorphic” here refers to the external notion of isomorphism in \mathbf{C} , involving definitional equality of maps; and so one cannot represent this definition internally in the type theory, since definitional equality is not represented by a type. Indeed, there is no way of defining these fibrations internally: every statement of the type theory respects equivalence, and we see in Lemma 2.2.11 below that every map is equivalent to a fibration.

Weak equivalences, by contrast, are defined first internally, as in Section 2.1 above:

Definition 2.2.4 (Voevodsky). A map $f: A \longrightarrow B$ is an *equivalence* if for each $b: B$ the homotopy fiber $\text{hfib}(f, b)$ is contractible.

(Note that this is simply a property of f , not extra structure, since being an equivalence is a proposition in the sense of Section 2.1.2.)

Take a map $f: A \longrightarrow B$ in \mathbf{C} to be in \mathcal{W} if “ $(\lambda x. f(x))$ is an equivalence” holds in \mathbf{C} .

With these definitions, we are now ready for the main theorem of the section:

Theorem 2.2.5. *\mathbf{C} , with \mathcal{W} and \mathcal{F} as described above, is a fibration category.*

We consider the various axioms in turn.

Lemma 2.2.6. *\mathcal{W} satisfies the 2-out-of-6 property.*

Proof. We first show the analogous statement internally.

Let f, g, h be composable maps, and suppose $f \cdot g$ and $g \cdot h$ are equivalences. Then:

- $(g \cdot f)^{-1} \cdot g \cdot (h \cdot g)^{-1}$ gives a quasi-inverse for $h \cdot g \cdot f$;
- $(h \cdot g)^{-1} \cdot h$ and $f \cdot (g \cdot f)^{-1}$ give left and right inverses for g ;
- $(g \cdot f)^{-1} \cdot g$ gives a quasi-inverse for f ;
- $g \cdot (h \cdot g)^{-1}$ gives a quasi-inverse for h .

This immediately implies the desired external statement, since internal and external composition agree. \square

Lemma 2.2.7. *Pullbacks of fibrations exist.*

Proof. The pullback of a dependent projection is given by substituting into the corresponding dependent type; that is, the following square is a pullback:

$$\begin{array}{ccc} \sum_{x:A'} B(fx) & \longrightarrow & \sum_{x:A} B(x) \\ \downarrow & & \downarrow \\ A' & \xrightarrow{f} & A. \end{array}$$

The two pullbacks lemma implies that pullbacks of their composites then also exist. \square

(Note that these really are pullbacks in the strict, external sense, in contrast to the internally defined pullbacks of Section 3.1, which from an external point of view become homotopy pullbacks.)

Lemma 2.2.8. *Given a fibration $\pi_1: \sum_{x:A} B(x) \longrightarrow A$, for any $a: A$ we have $B(a) \simeq \text{hfib}(\pi_1, a)$.*

Proof. Take any $a : A$. For the map $B(a) \longrightarrow \text{hfib}(\pi_1, a)$, send $b : B(a)$ to $((a, b), \text{refl}(a))$. Conversely, send $((a', b), p) : \text{hfib}(\pi_1, a)$ (where $b : B(a')$ and $p : (a' \rightsquigarrow a)$) to the transported element $p_!(b) : B(a)$. The verification that these are mutually inverse is straightforward. \square

Lemma 2.2.9. *Fibrations and acyclic fibrations are preserved under pull-back.*

Proof. Preservation of fibrations is clear by construction from the proof of Lemma 2.2.7; and for acyclic fibrations, combine this with Lemma 2.2.8 and the fact that equivalence preserves contractibility. \square

Definition 2.2.10. *Path objects* in \mathbf{C} are constructed from the identity types:

$$\begin{array}{ccc} & \text{PA} := \sum_{x, y: A} (x \rightsquigarrow y) & \\ & \nearrow & \searrow \\ A & \xrightarrow{\Delta} & A \times A \end{array}$$

We have now amassed all the ingredients of a fibration category:

Proof of Theorem 2.2.5. Immediate from the preceding lemmas. \square

Besides the basic structure of a fibration category, we consider a few more useful properties from the theory of fibration categories that are satisfied by \mathbf{C} :

Lemma 2.2.11 (Factorization Lemma, [GG08, Lem. 11]). *For every morphism $f : A \longrightarrow B$ in \mathbf{C} , there exists a factorization:*

$$\begin{array}{ccc} & \text{Pf} & \\ \sigma_f \nearrow & & \searrow p_f \\ A & \xrightarrow{f} & B \end{array}$$

with $\sigma_f \in \mathcal{W}$ and $p_f \in \mathcal{F}$.

Proof. We take

$$\text{Pf} := \sum_{y: B, x: A} (fx \rightsquigarrow y).$$

and

$$\sigma_f(x) := (fx, x, \text{refl}(fx)).$$

By definition, p_f is in \mathcal{F} ; and it is easy to check that $\sigma_f \in \mathcal{W}$. \square

Such $(\mathcal{W}, \mathcal{F})$ factorizations may be constructed in any fibration category. In the type-theoretic case, however, they crucially satisfy an additional property, corresponding to the Id -elimination rule: σ_f is weakly left-orthogonal to fibrations. We will not however go into this point here; see [GG08] for details.

Lemma 2.2.12 (Right properness). *The pullback of a weak equivalence along a fibration is again a weak equivalence:*

$$\begin{array}{ccc} \pi^*C & \xrightarrow{\pi^*f} & \sum_{x:A} B(x) \\ \downarrow & & \downarrow \pi \\ A' & \xrightarrow{f \in \mathcal{W}} & A \end{array}$$

Proof. The map π^*f sends a pair (y, b) to $(f(y), b)$; so taking a quasi-inverse (g, η, ϵ) for f , we can construct a quasi-inverse for π^*f by sending (x, b) to $(g(x), \eta(x)_! b)$. □

(Again, this property holds in any fibration category.)

One may also define cofibrancy, for objects of any fibration category:

Definition 2.2.13. An object C of a fibration category \mathbf{C} is *cofibrant* if for any acyclic fibration $p: B \rightarrow A$ and map $f: C \rightarrow A$, there is some lifting $\bar{f}: C \rightarrow B$:

$$\begin{array}{ccc} & & B \\ & \nearrow \bar{f} & \downarrow p \\ C & \xrightarrow{f} & A \end{array}$$

When \mathbf{C} is a categorical model of \mathcal{H}' (though not in general), we have:

Lemma 2.2.14. *Every object of \mathbf{C} is cofibrant.*

Proof. By Lemma 2.2.8 every acyclic fibration admits a section, so we may simply take \bar{f} to be the composite of this section with f . □

We conclude with a somewhat subtler question. Another condition often assumed for fibration categories is that any sequence

$$A_0 \xleftarrow{f_0} A_1 \xleftarrow{f_1} A_2 \xleftarrow{f_2} \dots,$$

in which each f_i is a fibration, has a limit, and that moreover the projections from this limit are again fibrations.

This turns out not to be provable in the type theory—in particular, it fails in the syntactic category. However, appropriate internally-formulated versions of it do hold; this is analogous to the fact that an elementary topos may fail to be externally complete, while possessing all limits in the internal sense.

To see how it fails in $\mathcal{C}(\mathcal{H}')$, consider the sequence of projections

$$1 \xleftarrow{f_0} \mathbf{Nat} \xleftarrow{f_1} \mathbf{Nat}^2 \xleftarrow{f_2} \dots$$

This sequence cannot have a limit, since such a limit would be a \mathbb{N} -fold product of copies of \mathbf{Nat} , and as such would necessarily have uncountably many global elements, while $\mathcal{C}(\mathcal{H}')$ is countable.

However, an *internal* limit for the sequence exists, in the form of the object $\mathbf{Nat}^{\mathbf{Nat}}$ (working internally, it does not make sense to ask whether the projections are fibrations); and, in some models (e.g. the simplicial model) this object turns out to be interpreted as the external limit $\prod_{\mathbb{N}} \mathbf{Nat}$.

3. LIMITS AND APPLICATIONS

Almost all the reasoning of this section takes place within the type theory; we therefore no longer distinguish such material explicitly.

3.1. Pullbacks and equalizers. Before defining general limits over graphs, we start by investigating pullbacks; these serve as both a warmup and a useful tool for subsequent material.

3.1.1. The standard construction of a pullback. We start by explicitly constructing the pullback of a cospan.

Definition 3.1.1. Let $A \xrightarrow{f} C \xleftarrow{g} B$ be a cospan of types and functions. The (*standard*) pullback $\mathbf{Pb}(f, g)$ of this cospan is defined as:

$$\mathbf{Pb}(f, g) := \sum_{x:A, y:B} (fx \rightsquigarrow gy)$$

with the obvious maps:

$$\begin{array}{ccc} \mathbf{Pb}(f, g) & \xrightarrow{\pi_B} & B \\ \pi_A \downarrow & & \downarrow g \\ A & \xrightarrow{f} & C \end{array}$$

(This definition may be recast to parallel a traditional construction of the homotopy pullback in fibration categories [Bro73, Lem 1.3]: first fibrantly replace f by p_f as in Lemma 2.2.11, obtaining

$$\begin{array}{ccc} & \mathbf{P}f = \sum_{c:C} \mathbf{hfib}(f, c) = \sum_{c:C, a:A} (fa \rightsquigarrow c) & \\ & \downarrow p_f & \\ A & \xrightarrow{f} & C, \end{array}$$

and then secondly, take the strict pullback of $\mathbf{P}f$ along g as a fibration over C , obtaining $\sum_{b:B} \mathbf{hfib}(f, g(c)) = \sum_{b:B, a:A} (fa \rightsquigarrow gb)$, which is (strictly, externally) isomorphic to $\mathbf{Pb}(f, g)$ as defined above.)

As a first basic property, note that the pullback is symmetric: there is an equivalence $\mathbf{Pb}(f, g) \simeq \mathbf{Pb}(g, f)$ commuting appropriately with the projections and canonical homotopy.

Moreover, the construction of the pullback is functorial in (f, g) ; this requires a few extra definitions to state:

Definition 3.1.2. Let $A \xrightarrow{f} C \xleftarrow{g} B$ and $A' \xrightarrow{f'} C' \xleftarrow{g'} B'$ be two cospans. A map h of cospans from (f, g) to (f', g') consists of maps h_A, h_B, h_C and homotopies h_f, h_g :

$$\begin{array}{ccccc}
 & & B & & \\
 & & \downarrow g & \searrow h_B & \\
 & & C & \xleftarrow{h_g} & B' \\
 A \xrightarrow{f} & & & & \downarrow g' \\
 & \searrow h_A & \nearrow h_f & \searrow h_C & \\
 & A' & \xrightarrow{f'} & C' &
 \end{array}$$

There is an identity map from any cospan to itself; also, there is an evident composition of cospan maps.

Proposition 3.1.3. *A map of cospans $h: (f, g) \longrightarrow (f', g')$ induces a map of pullbacks $\text{Pb}(f, g) \longrightarrow \text{Pb}(f', g')$. Moreover, this preserves composition and identities.*

The most frequent application of this functoriality, in practice, is the invariance of pullbacks under equivalences — that, for instance, given a cospan $A \xrightarrow{f} C \xleftarrow{g} B$ and an equivalence $e: A' \simeq A$, there is an equivalence between the pullbacks $\text{Pb}(f, g)$ and $\text{Pb}(f \cdot e, g)$. This, and various other similar statements, are all easily obtained from the functoriality of Pb together with the lemma:

Lemma 3.1.4. *Suppose $h = (h_A, h_B, h_C, h_f, h_g)$ is a cospan map from (f, g) to (f', g') , and h_A, h_B, h_C are equivalences. Then there is a cospan map $h^{-1}: (f', g') \longrightarrow (f, g)$, inverse to h in that there are paths $(h \cdot h^{-1} \rightsquigarrow 1)$ and $(h^{-1} \cdot h \rightsquigarrow 1)$.*

An interesting technical point arises here: rather than proving this and other facts about cospan maps directly, we deduce them from the analogous facts about commutative squares (considered as maps between functions). Most immediately, this slightly simplifies the proofs in the present section (since one does not have to write each construction out separately for the left and right legs of the cospan); but it also allows us to directly re-use the commutative squares material in Section 3.2, as the building blocks of the analogous facts about diagrams over general graphs.

3.1.2. The universal property of pullbacks. Alternatively, we can characterize pullbacks using a universal property. For the next few definitions, fix some cospan $A \xrightarrow{f} C \xleftarrow{g} B$.

Definition 3.1.5. Let X be any type. A *cone* μ over (f, g) with vertex X consists of functions μ_A, μ_B , and a homotopy $\mu_C : f \cdot \mu_A \Rightarrow g \cdot \mu_B$:

$$\begin{array}{ccc} X & \xrightarrow{\mu_B} & B \\ \mu_A \searrow & \nearrow \mu_C & \downarrow g \\ A & \xrightarrow{f} & C \end{array}$$

Write $\text{Cone}(X; f, g)$ for the type of cones over (f, g) with vertex X .

$\text{Cone}(X; f, g)$ should be contravariantly functorial in X . We do not show this in full; but in particular, a map $f : X' \rightarrow X$ induces a map

$$\text{Cone}(X; f, g) \rightarrow \text{Cone}(X'; f, g),$$

given by precomposing the components of the cone with f ; for a cone μ , we denote this as $\mu \circ f$. Fixing a cone $\mu : \text{Cone}(X; f, g)$ thus induces for any type X' a map

$$(\mu \circ -) : (X' \rightarrow X) \rightarrow \text{Cone}(X'; f, g).$$

This allows us to define the universal property of pullbacks:

Definition 3.1.6. A cone μ over (f, g) , with vertex P , is an (*abstract pullback*) for (f, g) if for every small type $X : \mathbf{U}$, the map $(\mu \circ -)$ gives an equivalence $(X \rightarrow P) \simeq \text{Cone}(X; f, g)$.

One can of course ask whether $(\mu \circ -)$ is an equivalence for an arbitrary type X , not necessarily small; but to quantify over types, one must restrict to some universe. Even doing so, the resulting property of “being a pullback” is (a priori) as large as the universe used. It is, however, an h-prop (since being an equivalence is one).

(For an investigation of left universal properties of inductive types, defined along similar lines, see [AGS12].)

Proposition 3.1.7. *The evident cone from the standard pullback $\text{Pb}(f, g)$ (3.1.1) to (f, g) is an abstract pullback.*

Proof. By direct construction: any cone from some X to (f, g) induces a map $X \rightarrow \text{Pb}(f, g)$, and by functional extensionality, this construction is inverse to composition with the standard cone. \square

Proposition 3.1.8. *If $\mu : \text{Cone}(X; f, g)$ and $\nu : \text{Cone}(Y; f, g)$ are both pullbacks for (f, g) , then the unique map $f : Y \rightarrow X$ such that $(\mu \circ f \rightsquigarrow \nu)$ (provided by the universal property of μ) is an equivalence.*

Conversely, if $\mu : \text{Cone}(X; f, g)$ is any cone, and $f : X \simeq Y$ an equivalence, then setting $\nu := \mu \circ f$, μ is a pullback if and only if ν is.

Proof. The following diagram commutes, and the maps $X \longrightarrow (1 \rightarrow X)$, $Y \longrightarrow (1 \rightarrow Y)$ are equivalences:

$$\begin{array}{ccc}
 Y & \longrightarrow & (1 \rightarrow Y) \\
 f \downarrow & & \searrow (\nu \circ -) \\
 X & \longrightarrow & (1 \rightarrow X) \xrightarrow{(\mu \circ -)} \text{Cone}(1; f, g)
 \end{array}$$

It follows by 2-out-of-3 that if any two of f , $(\mu \circ -)$, $(\nu \circ -)$ are equivalences, so is the third. \square

Corollary 3.1.9. *A cone $\mu : \text{Cone}(X; f, g)$ is a pullback if and only if the induced map $X \longrightarrow \text{Pb}(f, g)$ is an equivalence.*

Since any two interderivable h-props are necessarily equivalent, this property could be used as an alternative definition of μ being a pullback, with the advantage of being again a small type.

3.1.3. Two pullbacks lemmas.

Proposition 3.1.10 (Concrete two pullbacks lemma). *For all f, g, h as in the diagrams below, the map $\text{Pb}(g^*f, h) \longrightarrow \text{Pb}(f, g \cdot h)$ induced by the composition of the two squares is an equivalence:*

$$\begin{array}{ccccc}
 & & \text{Pb}(f, g \cdot h) & & \\
 & \nearrow & & \searrow & \\
 & \text{Pb}(g^*f, h) & \longrightarrow & \text{Pb}(f, g) & \longrightarrow A \\
 & \downarrow & & \downarrow g^*f & \downarrow f \\
 B' & \xrightarrow{h} & B & \xrightarrow{g} & C
 \end{array}$$

Proposition 3.1.11 (Abstract two pullbacks lemma). *Suppose that in a rectangle*

$$\begin{array}{ccccc}
 P' & \longrightarrow & P & \longrightarrow & A \\
 \downarrow & & \downarrow k & & \downarrow f \\
 B' & \xrightarrow{h} & B & \xrightarrow{g} & C
 \end{array}$$

the lower square is a pullback. Then the upper square is a pullback if and only if the outer rectangle is a pullback.

Proof. Write μ for the cone from P to (f, g) ; μ' for the cone from P' to (g^*f, h) ; and μ'' for the cone from P' to $(f, g \cdot h)$. Then for any X , the

following triangle commutes:

$$\begin{array}{ccc}
 & & \text{Cone}(X; g^* f, h) \\
 & \nearrow^{(\mu' \cdot -)} & \downarrow \\
 (X \rightarrow P') & & \\
 & \searrow_{(\mu'' \circ -)} & \text{Cone}(X; f, g \cdot h)
 \end{array}$$

Here the vertical map denotes the composition of a cone on $(g^* f, h)$ with μ ; and this can be shown (by direct construction) to be an equivalence. Hence by 2-out-of-3, $(\mu' \circ -)$ is an equivalence if and only if $(\mu'' \circ -)$ is. \square

It should be noted that the arguments involved in showing the equivalence $\text{Cone}(X; g^* f, h) \simeq \text{Cone}(X; f, g \cdot h)$ are necessarily more involved than in the 1-categorical setting, since they depend on comparing paths in types; in terms of the classical theory, this is more analogous to the corresponding lemma for quasi-pullbacks in a bicategory.

3.1.4. Equalizers.

Definition 3.1.12. Let $f, g: A \longrightarrow B$. The *equalizer* of f and g is defined as the type:

$$\text{Eq}(f, g) := \sum_{x:A} (fx \rightsquigarrow gx).$$

together with the projection $\pi: \text{Eq}(f, g) \longrightarrow A$.

As in classical category theory, pullbacks and equalizers can be defined in terms of each other.

Proposition 3.1.13. *The equalizer of any pair $f, g: A \longrightarrow B$ is equivalent to the pullback of the paired map $\langle f, g \rangle: A \longrightarrow B \times B$ and the diagonal Δ_B :*

$$\begin{array}{ccc}
 \text{Eq}(f, g) \simeq \text{Pb}(\Delta_B, \langle f, g \rangle) & \longrightarrow & A \\
 \downarrow & & \downarrow \langle f, g \rangle \\
 B & \xrightarrow{\Delta_B} & B \times B
 \end{array}$$

Conversely, the pullback of any cospan $A \xrightarrow{f} C \xleftarrow{g} B$ is equivalent to the equalizer of the pair

$$f \cdot \pi_1, g \cdot \pi_2: A \times B \longrightarrow C.$$

3.1.5. Homotopy fibers and loop spaces. We next consider a couple of examples which bring out the homotopical character of the theory—examples which in classical 1-category theory, and in the type theory with UIP, would be trivial, but which in the un-truncated type theory become non-trivial, corresponding to the classical theory of *homotopy* pullbacks.

For these we first need one piece of notation. Given a type B and an element $b : B$, write $\ulcorner b \urcorner : 1 \longrightarrow B$ for the map sending the unique element $*$: 1 to b .

Example 3.1.14. Given a map $f : A \longrightarrow B$ and an element $b : B$, the homotopy fiber of f over b may equivalently be given as a pullback:

$$\begin{array}{ccc} \text{hfib}(f, b) \simeq \text{Pb}(\ulcorner b \urcorner, f) & \longrightarrow & A \\ \downarrow & & \downarrow f \\ 1 & \xrightarrow{\ulcorner b \urcorner} & B \end{array}$$

Example 3.1.15. Given a type B and an element $b : B$, the space of loops in B based at b , $\Omega(B, b) := (b \rightsquigarrow_B b)$ may be given as a pullback:

$$\begin{array}{ccc} \Omega(B, b) \simeq \text{Pb}(\ulcorner b \urcorner, \ulcorner b \urcorner) & \longrightarrow & 1 \\ \downarrow & & \downarrow \ulcorner b \urcorner \\ 1 & \xrightarrow{\ulcorner b \urcorner} & B. \end{array}$$

3.1.6. *Properties of pullbacks.* Various nice properties of maps are preserved under pullback. In proving such preservation properties, the following lemma is rather useful:

Proposition 3.1.16. *Given $A \xrightarrow{f} C \xleftarrow{g} B$, the homotopy fiber of the f^*g over a point $a : A$ is equivalent to the homotopy fiber of g over $f(a)$.*

Proof. This follows from Example 3.1.14 combined with the two pullbacks lemma. \square

From this, we can immediately obtain:

Corollary 3.1.17. *Let $A \xrightarrow{f} C \xleftarrow{g} B$ be a cospan. If f is an equivalence, then so is the projection π_B from the pullback.*

Similarly, any property of maps that can be characterized fiberwise is stable under pullback.

3.2. **Limits.** Generalizing the constructions of pullbacks and equalizers above, we move to limits for diagrams over arbitrary graphs.

3.2.1. *Graphs and diagrams.*

Definition 3.2.1. A *graph* G consists of:

- a type G_0 (the *vertices* or *objects* of G); and
- for each $i, j : G_0$, a type $G_1(i, j)$ (the *edges* or *arrows* from i to j).

Definition 3.2.2. A *diagram* D on a graph G consists of:

- for each vertex $i : G_0$, a type $D_0(i)$;
- for each arrow $g : G_1(i, j)$ of G , a map $D_1(g) : D(i) \longrightarrow D(j)$.

For both graphs and diagrams, we will often suppress the subscripts when they are clear from context.

Example 3.2.3. To recover cospans as an example of these diagrams, one can define a graph by taking G_0 to be the type with three elements, $\{l, m, r\}$ and let G_1 be given by:

- $G(l, m) := 1$,
- $G(r, m) := 1$,
- $G(i, j) := \emptyset$ otherwise.

A diagram D over this graph corresponds precisely to a cospan:

$$\begin{array}{ccc} & D(r) & \\ & \downarrow & \\ D(l) & \longrightarrow & D(m) \end{array}$$

3.2.2. The universal property of limits.

Definition 3.2.4. Given a diagram D on a graph G , a *cone* μ on D with vertex X consists of:

- for each $i : G_0$, a map $\mu_i^0 : X \longrightarrow D_0(i)$;
- for each arrow $g : G_1(i, j)$, a homotopy $\mu_g^1 : D_1(g) \cdot \mu_i^0 \Rightarrow \mu_j^0$.

Write $\text{Cone}(X; D)$ for the type of cones on D with vertex X .

Again, we usually suppress the subscripts, writing just μ_i, μ_f .

As with cones over cospans, $\text{Cone}(X; D)$ is functorial in X ; a map $f : X' \longrightarrow X$ and a cone $\mu : \text{Cone}(X; D)$ may be composed to give a cone $\mu \circ f : \text{Cone}(X'; D)$. This now lets us generalize the definition of the universal property:

Definition 3.2.5. Let D be a diagram on the graph G . A cone μ over D , with vertex L , is an *(abstract) limit* for D if for every small type $X : \mathbf{U}$, the map $(\mu \circ -) : (X \rightarrow L) \longrightarrow \text{Cone}(X; D)$ is an equivalence.

By abuse of notation, we often speak of L being the limit of D , when the cone μ is implicit.

Most of the theorems of the preceding section generalize immediately. In particular,

Proposition 3.2.6. *Given any two limit cones for the same diagram, the canonical map between their vertices is an equivalence; conversely, the composition of any limit cone with an equivalence is again a limit cone.*

Again as in the previous section, there is a standard construction of the limit:

Definition 3.2.7. Let D be a diagram over a graph G . The *(standard) limit* $\text{Lim}D$ is the type of pairs (x, α) , where

- $x : \prod_{i : G_0} D(i)$;
- $\alpha : \prod_{i, j : G_0, g : G(i, j)} ((D(g)(x_i) \rightsquigarrow x_j))$.

There is an evident cone from $\text{Lim}D$ to D , and as one would hope,

Proposition 3.2.8. *$\text{Lim}(D)$ is an abstract limit for D .*

Proposition 3.2.9. *A cone μ from X to some diagram D is a limit for D if and only if the map $X \longrightarrow \text{Lim}D$ induced by μ is an equivalence.*

One again, we may define maps of diagrams, and show that Lim is functorial in such maps, and in particular, is functorial in equivalences. Since graphs, diagrams, and limits are all simply built up from arrows, these definitions and results follow straightforwardly once one has given the basic case of commutative squares, seen as maps between functions.

3.2.3. Examples and properties.

Example 3.2.10. In Example 3.2.3, we saw that cospans correspond to diagrams over a certain graph. Then cones over these diagrams correspond to cones over the cospans, as originally defined; and a diagram-cone is a limit exactly if the corresponding cospan-cone is a pullback.

Example 3.2.11. Just as in the classical 1-categorical theory, the limit over a diagram D may be constructed as an equalizer of maps between products:

$$\prod_{i,j:G, g:G(i,j)} D(i) \rightrightarrows \prod_{i:G} D(i)$$

Various useful facts are also straightforward to deduce from the standard construction; for instance,

Proposition 3.2.12. *If D is a diagram on some graph, and each type $D(i)$ has h -level n , then $\text{Lim}D$ has h -level n ; hence via the canonical equivalence, so does any other limit for D .*

3.2.4. Why not categories? One might reasonably ask here: why have we considered limits only over graphs, not over categories as is usual in the 1-categorical theory?

The problem—as ever in homotopical settings—is one of *coherence*. Defining a category internally is roughly analogous to defining an $(\infty, 1)$ -category externally; that is, it requires not only identity, composition, associativity, and the like, but also higher-dimensional data ensuring the coherence of the paths witnessing the associativity axioms, and so on in arbitrarily high dimensions. While we hope that this will eventually be possible in the type theory, it is currently far from clear how to present it.

In defining categories, this problem can be avoided by putting h -level restrictions on the morphisms; see [AKS13] for a development of the resulting theory. However, to talk about diagrams of arbitrary types over such categories would once again require an infinite family of coherence conditions, essentially since one is presenting an ∞ -functor into the $(\infty, 1)$ -category of all types, which is not generally of any finite h -level.

However, working with graphs avoids these issues entirely: a map out of a graph (or equivalently, out of the free category thereon) consists purely of 0- and 1-dimensional data, with no coherence required. (More generally, one could use a similar approach to describe diagrams over finite-dimensional computads, or semi-simplicial objects, without confronting coherence issues.)

3.3. Pointed types and fiber sequences.

3.3.1. Definitions.

Definition 3.3.1. A *pointed type* (A, a_0) is a type A , together with an element $a_0 : A$, the *basepoint*. (We will often refer to the both pointed type and its underlying type as A , and write $\text{pt}(A)$ for the basepoint.)

A *map of pointed types* (or *pointed map*) $(f, p) : (A, a_0) \longrightarrow (B, b_0)$ consists of a function $f : A \longrightarrow B$, together with a path $p : (f(a_0) \rightsquigarrow_B b_0)$. (Again, we will often write f for the whole pointed map, and write $\text{pt}(f)$ for its associated path.)

The loop space construction Ω lifts naturally to a map from pointed types to pointed types, setting $\Omega A := ((\text{pt } A \rightsquigarrow \text{pt } A), \text{refl}(\text{pt } A))$. One can therefore iterate it, giving the n -fold loop spaces $\Omega^n A$ of a pointed type. Moreover, this has an associated action on maps. A pointed map $f : A \longrightarrow B$ induces a pointed map $\Omega(f) : \Omega A \longrightarrow \Omega B$, with underlying map sending $q : (\text{pt } A \rightsquigarrow \text{pt } A)$ to $\overline{\text{pt } f} \cdot f[q] \cdot \text{pt } f : (\text{pt } B \rightsquigarrow \text{pt } B)$.

Similarly, the homotopy fiber construction hfib lifts naturally to the pointed world. Given a pointed map $f : A \longrightarrow B$, write $\text{hfib}(f)$ for the pointed type given by $\text{hfib}(f, \text{pt } B)$, with basepoint $(\text{pt } A, \text{pt } f)$; and the inclusion $\text{hfib}(f) \longrightarrow A$ is again a pointed map.

3.3.2. The long exact sequence theorem. As an application of the above tools, we can now recover the long exact sequence associated to a pointed map.

Definition 3.3.2. A *fiber sequence* is a pair of pointed maps $F \xrightarrow{g} E \xrightarrow{f} B$, together with an equivalence $F \simeq \text{hfib}(f)$ commuting with the inclusion $\text{hfib}(f) \longrightarrow E$.

Note that up to canonical equivalence, a fiber sequence is determined simply by the single pointed map $E \longrightarrow B$.

Theorem 3.3.3. *Given a pointed map $f : E \longrightarrow B$, there is a sequence of maps:*

$$\dots \longrightarrow \Omega^2 B \longrightarrow \Omega F \longrightarrow \Omega E \longrightarrow \Omega B \longrightarrow F \longrightarrow E \longrightarrow B$$

in which every pair of consecutive maps forms a fiber sequence.

Proof. Taking $F := \text{hfib}(f)$, it is sufficient to prove that the homotopy fiber of the inclusion $F \longrightarrow E$ is equivalent to ΩB ; subsequent stages follow by iteration. One can prove this equivalence by direct construction; alternatively, the results of Section 3.1 allow us to give a rather more conceptual proof, due originally to Mather [Mat76, Lem. 32]:

$$\begin{array}{ccccc} \bullet & \longrightarrow & F & \longrightarrow & 1 \\ \downarrow & & \downarrow & & \downarrow \lceil \text{pt } B \rceil \\ 1 & \xrightarrow{\lceil \text{pt } E \rceil} & E & \xrightarrow{f} & B \end{array}$$

By the two pullbacks lemma, the pullback of the left-hand square is equivalent to the pullback of the whole rectangle. But by Examples 3.1.14 and 3.1.15, these pullbacks are respectively equivalent to the homotopy fiber of $F \longrightarrow E$, and to the loop space ΩB . \square

4. REFLECTIONS ON THE FORMAL VERIFICATION

Formalizing the constructions Sections 2 and 3 was often straightforward: many of the definitions are very naturally expressed in the language of type theory, and verifying their properties is often just a matter of unpacking definition and applying straightforward logical manipulations and background facts. Sometimes, however, additional effort was required. In this section, we survey some of the practical lessons learned during the formalization.

4.1. Limitations. One fundamental challenge that arises comes from working purely in the type theory. In classical approaches to homotopy theory, one always has an extra external scaffolding available, with (in particular) strict, on-the-nose equality on all types of objects. One typically expects the main results and constructions to respect appropriate notions of equivalence, but one is free to use intermediate constructions that do not.

Developing the homotopy theory in HoTT, we are constrained to work entirely in a homotopy-invariant manner, rendering some classical techniques unavailable. In most cases, some fully invariant approach is reasonably apparent; but sometimes, one is not. We saw such a case in Section 3.2: we do not know how to represent the notion of a diagram over an arbitrary category, and so restricted attention to (diagrams and limits over) graphs.

4.2. Proof-relevance. Another difficulty lies in getting used to thinking of proofs of equalities as *constructions* that one might need to prove things about later on.

In traditional formalizations, equalities are *proof-irrelevant*: different proofs of the same equality are not logically distinguishable. In Coq, for instance, one could safely end them with the keyword `Qed`, which renders them *opaque*, meaning that one cannot later access their contents. In traditional mathematics, this makes sense; once one has an equality, one only needs the fact that it holds, treating the proof as a black box.

In HoTT, however, equality is proof-relevant: a path type may have multiple logically distinct inhabitants. When constructing equality proofs in this setting, one typically needs to end an equality proof with the keyword `Defined`, allowing the user to unfold that definition later on. The specifics of the proof matter; one tries to keep proofs as clean and short as possible, using lemmas and constructions with known, previously proven properties. Unfortunately, this means that several of Coq’s powerful tactics (notably the `rewrite` family) are somewhat unsatisfactory in our setting: the paths they produce are difficult to reason about later.

On the other hand, some important statements remain proof-irrelevant. If a type has been shown to be an h-prop, one knows that any two elements of it are canonically equal; so one may make such an element opaque without losing any logical content. Even so, it is often convenient to leave such objects transparent, to retain their *computational* content.

For instance, for a function `f`, the type `is_equiv f` (the property that `f` is an equivalence) is an h-prop; so in principle one may safely render a proof of this opaque. However, one often uses such a proof to produce an inverse for `f`; if the proof was transparent, then the resulting inverse will retain computational properties from its construction, whereas if the proof is opaque, one must reason explicitly about the action of the inverse. We formed no clear convention on this: sometimes it turned out more convenient to keep such proofs transparent, for easier reduction in later proofs; in other case, this was unnecessary, and making the proofs opaque gave more efficient compilation.

4.3. Constructing paths. The most fundamental type constructor in homotopy type theory is the type of paths, and the most challenging parts of proofs usually involved constructing paths between complex objects. Given the subject matter, we never had to pass beyond the 2-categorical level, constructing paths between paths; but even so, this required a good deal of care, and facility with path algebra.

One recurring situation was the construction of paths between elements of a dependent sum, or elements of a record type with dependencies between components. For example, if (a, b) and (a', b') are elements of a type $\sum_{x:A} B(x)$, constructing a path between these two elements involves constructing a path p from a to a' , and then constructing a path q from the transport of b along p to b' . Thus in general we have:

```

Lemma total_paths {A : Type} {B : A -> Type}
  {s s' : total B}
  (p : paths (pr1 s) (pr1 s'))
  (q : paths (transport p (pr2 s)) (pr2 s'))
  : paths s s'.

```


where `pr1` and `pr2` denote the two projections from the total space $\sum_{x:A} B(x)$. For interactive, tactic-based proofs, we generally found it useful to bundle the arguments p, q into a single structure:

```

Lemma total_paths' {A : Type} {B : X -> Type}
  {s s' : total B}
  : { p : pr1 s = pr1 s' & p # pr2 s = pr2 s' } -> s = s'.

```

Recall that here `{ p : pr1 s = pr1 s' & p # pr2 s = pr2 s' }` is notation for a dependent sum, denoting the type of pairs (p, q) as above. When constructing a path between elements of a dependent sum, even when p is explicitly available, applying `(total_paths p)` sometimes fails to infer implicit arguments. Instead, applying `total_paths'` leaves the goal of providing the pair (p, q) , providing the user explicitly with their required types. The tactic `exists p` can then be used to give the first component, leaving the goal of constructing the second path q interactively.

The problem is that *transport* is rather difficult to work with. There are many library lemmas about how its behaviour depends on the dependent type B , which in principle allow one to work with transported terms; but we found it more convenient to directly give tailored variants of `total_paths` for each specific Σ - and record type.

For example, taking a cospan $f: A \rightarrow C, g: B \rightarrow C$, the standard pullback of f and g is given by the type $\sum_{x:A, y:B} (fx \rightsquigarrow gy)$. Using `total_path` to provide a path in this type between triples $(x; (y; p))$, $(x'; (y'; p'))$ would require three paths $q : x = x'$, $r : q \# y = y'$, and $s : r \# q \# p = p'$. Notice, however, that in this case the second component, y , does not depend on x , so the transport is trivial; and moreover, the doubly-transported third component can be explicitly described as a composite. Thus, one can provide the following lemma to construct a path between two elements of the standard pullback:

```

Definition pullback_path' {A B C : Type} {f : A -> C}
  {g : B -> C} (u u' : pullback f g)
  : { q : pullback_proj1 u = pullback_proj1 u'
    & {r : pullback_proj2 u = pullback_proj2 u'
    & !(map f q) @ (pullback_comm u) @ (map g r)
    = pullback_comm u' } }
  -> u = u'.

```

The process of analyzing the canonical data for presenting a path between elements of a complex type, and writing lemmas to construct and work with such paths, was crucial to the formalization.

To consider one last example of this sort, recall that a cospan cone, that is, a diagram on the data f, g above, consists of a space, X , and maps h and k from X to A and B , respectively, making the diagram commute.

```

Definition cospan_cone {A B C : Type} (f : A -> C)

```

```

(g : B -> C) (X : Type)
:= { h : (X -> A) &
    { k : (X -> B) & forall x, paths (f(h x)) (g(k x)) }}.

```

A path between two such cones involves, in particular, a path between the family of paths in the third component:

Definition `cospan_cone_path`

```

{A B C : Type} {f : A -> C} {g : B -> C} {X : Type}
{Phi1 Phi2 : cospan_cone f g X}
(p : cospan_cone_map1 Phi1 = cospan_cone_map1 Phi2)
(q : cospan_cone_map2 Phi1 = cospan_cone_map2 Phi2)
(r : forall x:X,
  cospan_cone_comm Phi1 x
  = (map f (happly p x))
    @ cospan_cone_comm Phi2 x
    @ !(map g (happly q x)))
: Phi1 = Phi2.

```

Here, `cospan_cone_map1`, `cospan_cone_map2`, and `cospan_cone_comm` refer to the three components of a cospan cone in the preceding definition. As with `total_paths`, we also give a version `cospan_cone_path'` that packages the required components into a dependent sum, and is often more convenient in interactive proofs.

The advantage to these formulations is that is comparatively straightforward (using lemmas from the HoTT library) to reason about operations like `map`, `happly`, `funext`, and their interactions with each other, as well as with path operations such as concatenation and inversion.

Returning to the question of this path-algebra itself, we found the formalization to require significant facility with such calculations. The HoTT library has a number of tactics for automating common manipulations and simplifications, but we found these tactics generally slowed down the proof-checker significantly. So, for the most part, we ended up giving such calculations by hand, building them explicitly from basic lemmas.

4.4. General strategies. We found it important to develop our theories and proofs in a modular way. The value of modularity in interactive theorem proving is well understood (see, for example, [G⁺13]), but in the context of homotopy type theory, it takes on additional significance. For one thing, many statements involving paths can only be proved when stated in full generality (to make available the elimination for `Id-types`). As a consequence, some facts cannot be derived in the course of a proof, on the fly, but have to be expressed independently. The fact that one often needs to reason about the construction of paths provides an additional reason to construct such proofs out of individually-named component lemmas: doing so allows one

derive properties of the components individually, and then invoke these properties later on. In other words, reasoning about a modularly-constructed proof allows one to work with the individual lemmas and unpack their contents selectively, as needed. In contrast, the failure to modularize can result in formal terms that are overwhelming in complexity.

Perhaps the most important lesson we learned was not to expect too much from an interactive theorem prover. Although homotopy type theory provides a powerful framework to support homotopy-theoretic reasoning, one still needs a thorough understanding of the relevant mathematics. To get some of the more complex proofs and constructions to work, we found it vitally important to find the right definitions, the right way of formulating assertions, the right supporting infrastructure, and the right proof strategies. This required thinking carefully about the mathematical content, avoiding the temptation to simply dive in and hack.

This should not suggest that Coq was no help at all. Indeed, Coq was excellent for helping us keep track of definitions and formulate statements correctly. Especially for more complex path-constructions, applying standard rules to unwrap and reduce the contents of a goal type was an extremely useful aid to finding the term required. In practice, we found ourselves going back and forth between the blackboard and Coq, using Coq to negotiate the inevitable syntactic bureaucracy, and then returning to the blackboard to recoup intuitions and plan proof-strategies. In this way, Coq earned its keep, serving as a “proof assistant” in a very real sense.

4.5. A case study: the two pullbacks lemma. We close with a discussion of the abstract two pullbacks lemma, Proposition 3.1.11, by way of illustration. Somewhat to our surprise, this turned out to be the most difficult proof in our formalization. In the end, we tried three substantially different approaches before finding one satisfactory.

Consider for now just the forward direction of Proposition 3.1.11, which states that if both squares have the universal property of pullbacks, then so does the composite. Let $f: A \rightarrow C$, $g: B \rightarrow C$, $h: B' \rightarrow C$, and $k: P \rightarrow B$ denote the maps so labeled in the diagram there. Our first approach invoked the concrete two pullbacks lemma, Proposition 3.1.10, which states that

$$\text{Pb}(g^*f, h) \simeq \text{Pb}(f, g \cdot h).$$

We then derived the following chain of equivalences, using the fact that cones from X to the cospan (f, g) are equivalent to maps from X to the standard pullback:

$$\begin{aligned} (X \rightarrow P_2) &\simeq \text{Cone}(X; k, h) \\ &\simeq (X \rightarrow \text{Pb}(k, h)) \\ &\simeq (X \rightarrow \text{Pb}(g^*f, h)) \\ &\simeq (X \rightarrow \text{Pb}(f, g \cdot h)) \\ &\simeq \text{Cone}(X; f, g \cdot h). \end{aligned}$$

Here the second and last equivalences are just the universal properties of the concrete pullbacks. The notation g^*f in the third equivalence denotes the pullback of f along g according to the concrete pullback construction; this equivalence relies on the fact that any abstract pullback is equivalent to the concrete one, and the fact that the concrete pullback construction is functorial. The equivalence of the left- and right-hand sides of the chain above *almost* gives what we want: however, the universal property for the outer pullback square requires not just that an equivalence exists, but that the *canonical* map from $X \longrightarrow P_2$ to $\mathbf{Cone}(X; f, g \cdot h)$ is an equivalence.

What remains is thus to show that the map we have just constructed is homotopic to the canonical one! This, however, turned out to be extremely difficult. The problem was a failure of modularity: all we could do was unwrap the long, complicated term, and calculate. We managed to do this, but although the tactic engine declared the effort successful, we were unable to get it past the type-checker (presumably because the resulting term was too large).

Our second approach involved constructing the desired inverse by hand. Any cone $\mu : \mathbf{Cone}(X; f, g \cdot h)$ over the outer cospan can be reinterpreted as a cone $\mu' : \mathbf{Cone}(X; f, g)$ over the bottom cospan. Applying the universal property of the cone from P_1 , we obtain a map $m_1 : X \longrightarrow P_1$ inducing μ' ; we can then take m_1 as the first leg of a cone $\mu'' : \mathbf{Cone}(X; k, h)$ on the top cospan. Applying the universal property of the cone on P_2 then gives a map $m_2 : X \longrightarrow P_2$, as desired. However, the task of proving that this construction is indeed a two-sided inverse for $(\mu \circ -)$ turned out to be difficult. For example, the first task requires one to show that, starting with a cone $\mu : \mathbf{Cone}(X; f, g \cdot h)$, carrying out the procedure above to obtain a map from X to P_2 and then taking the induced cone, the resulting cone $\nu : \mathbf{Cone}(X; f, g)$ is connected by a path to the original μ . As described in Section 4.3, this involves showing not only that the component maps agree, but also that the resulting families of equality proofs agree as well; this turns out to be an interesting but laborious exercise in bicategorical path-algebra.

We finally settled on the approach described in Section 3.1.3, which establishes both directions of Proposition 3.1.11 simultaneously. Showing that the type $\mathbf{Cone}(X; k, h)$ of cones on the top cospan is equivalent to the type $\mathbf{Cone}(X; f, g \cdot h)$ of cones on the outer cospan required some effort, but the result was still considerably cleaner than either of the previous proofs. With that in hand, all that remained was to show that the triangle depicted in the proof of Proposition 3.1.11 in Section 3.1.3 commutes. To our very pleasant surprise, this fact had a one-line proof in Coq:

```

Lemma two_pullback_triangle_commutates {P1 : Type}
  (C1 : cospan_cone f g P1)
  {P2 : Type} (C2 : cospan_cone (cospan_cone_map2 C1) h P2)
  {X : Type} (m : X -> P2)
: top_cospan_cone_to_composite C1 (map_to_cospan_cone C2 X m)

```

```
= map_to_cospan_cone (top_cospan_cone_to_composite C1 C2) X m.
```

```
Proof.
```

```
  exact (idpath _).
```

```
Defined.
```

In other words, the left- and right-hand sides are definitionally equal.

REFERENCES

- [AGS12] Steve Awodey, Nicola Gambino, and Kristina Sojakova, *Inductive types in homotopy type theory*, 2012 27th Annual IEEE Symposium on Logic in Computer Science (LICS), IEEE, 2012, pp. 95–104.
- [AKS13] Benedikt Ahrens, Krzysztof Kapulkin, and Michael Shulman, *Univalent categories and Rezk completion*, submitted, 2013.
- [AW09] Steve Awodey and Michael A. Warren, *Homotopy theoretic models of identity types*, Math. Proc. Cambridge Philos. Soc. **146** (2009), no. 1, 45–55, [arXiv:0709.0248](#), [doi:10.1017/S0305004108001783](#).
- [Bau89] Hans Joachim Baues, *Algebraic homotopy*, Cambridge Studies in Advanced Mathematics, vol. 15, Cambridge University Press, Cambridge, 1989, [doi:10.1017/CB09780511662522](#).
- [Bro73] Kenneth S. Brown, *Abstract homotopy theory and generalized sheaf cohomology*, Trans. Amer. Math. Soc. **186** (1973), 419–458.
- [G⁺13] Georges Gonthier et al., *A machine checked proof of the odd order theorem*, 2013, preprint.
- [GG08] Nicola Gambino and Richard Garner, *The identity type weak factorisation system*, Theoret. Comput. Sci. **409** (2008), no. 1, 94–109, [arXiv:0803.4349](#), [doi:10.1016/j.tcs.2008.08.030](#).
- [HoT] HoTT group, *Homotopy type theory repository*, ongoing Coq development, <https://github.com/HoTT/coq>.
- [HS98] Martin Hofmann and Thomas Streicher, *The groupoid interpretation of type theory*, Twenty-five years of constructive type theory (Venice, 1995), Oxford Logic Guides, vol. 36, Oxford Univ. Press, New York, 1998, pp. 83–111.
- [KLV12] Krzysztof Kapulkin, Peter LeFanu Lumsdaine, and Vladimir Voevodsky, *The simplicial model of univalent foundations*, preprint, 2012, [arXiv:1211.2851](#).
- [Lum09] Peter LeFanu Lumsdaine, *Weak ω -categories from intensional type theory (conference version)*, Typed lambda calculi and applications (Berlin), Lecture Notes in Computer Science, vol. 5608, Springer, 2009, pp. 172–187.
- [Lum11] ———, *Model structures from higher inductive types*, unpublished note, December 2011, <http://www.mathstat.dal.ca/~p.l.lumsdaine/research/Lumsdaine-Model-strux-from-HITs.pdf>.
- [Lur09] Jacob Lurie, *Higher topos theory*, Annals of Mathematics Studies, vol. 170, Princeton University Press, Princeton, NJ, 2009.
- [Mat76] Michael Mather, *Pull-backs in homotopy theory*, Can. J. Math **28** (1976), no. 2, 225–263.
- [ML84] Per Martin-Löf, *Intuitionistic type theory*, Studies in Proof Theory. Lecture Notes, vol. 1, Bibliopolis, Naples, 1984.
- [PW12] Álvaro Pelayo and Michael Warren, *Homotopy type theory and Voevodsky’s univalent foundations*, preprint, 2012, [arXiv:1210.5658](#).
- [RB06] Andrei Radulescu-Banu, *Cofibrations in homotopy theory*, preprint, 2006, [arXiv:0610009](#).
- [RS13] Egbert Rijke and Bas Spitters, *Limits and colimits in Homotopy Type Theory*, in preparation, 2013.

- [Shu12] Michael Shulman, *The univalence axiom for inverse diagrams*, preprint, 2012, [arXiv:1203.3253](#).
- [Str91] Thomas Streicher, *Semantics of type theory*, Progress in Theoretical Computer Science, Birkhäuser Boston Inc., Boston, MA, 1991, Correctness, completeness and independence results, With a foreword by Martin Wirsing.
- [vdBG11] Benno van den Berg and Richard Garner, *Types are weak ω -groupoids*, Proc. Lond. Math. Soc. (3) **102** (2011), no. 2, 370–394, [arXiv:0812.0298](#), [doi:10.1112/plms/pdq026](#).
- [vdBG12] ———, *Topological and simplicial models of identity types*, ACM Trans. Comput. Log. **13** (2012), no. 1, Art. 3, 44, [arXiv:1007.4638v1](#), [doi:10.1145/2071368.2071371](#).
- [Voe] Vladimir Voevodsky, *Univalent foundations repository*, ongoing Coq development, <https://github.com/vladimirias/Foundations>.
- [Voe06] ———, *A very short note on homotopy λ -calculus*, notes from seminars given at Stanford University, 2006, http://math.ucr.edu/home/baez/Voevodsky_note.ps.
- [Voe10] ———, *Notes on type systems*, ongoing unpublished manuscript, 2010, http://www.math.ias.edu/~vladimir/Site3/Univalent_Foundations_files/expressions_current.pdf.

(Jeremy Avigad) CARNEGIE MELLON UNIVERSITY, PITTSBURGH
E-mail address: avigad@cmu.edu

(Krzysztof Kapulkin) INSTITUTE FOR ADVANCED STUDY, PRINCETON; AND UNIVERSITY OF PITTSBURGH
E-mail address: krk56@pitt.edu

(Peter LeFanu Lumsdaine) INSTITUTE FOR ADVANCED STUDY, PRINCETON
E-mail address: plumsdaine@ias.edu